

Succinct Approximate Rank Queries

Ran Ben Basat

Department of Computer Science, Technion

srn@cs.technion.ac.il

Abstract

We consider the problem of summarizing a multi set of elements in $\{1, 2, \dots, n\}$ under the constraint that no element appears more than ℓ times. The goal is then to answer *rank* queries — given $i \in \{1, 2, \dots, n\}$, how many elements in the multi set are smaller than i ? — with an additive error of at most Δ and in constant time. For this problem, we prove a lower bound of $\mathcal{B}_{\ell, n, \Delta} \triangleq \left\lfloor \frac{n}{\lceil \Delta / \ell \rceil} \right\rfloor \log(\max\{\lfloor \ell / \Delta \rfloor, 1\} + 1)$ bits and provide a *succinct* construction that uses $\mathcal{B}_{\ell, n, \Delta}(1 + o(1))$ bits. Next, we generalize our data structure to support processing of a stream of integers in $\{0, 1, \dots, \ell\}$, where upon a query for some $i \leq n$ we provide a Δ -additive approximation for the sum of the *last* i elements. We show that this too can be done using $\mathcal{B}_{\ell, n, \Delta}(1 + o(1))$ bits and in constant time. This yields the first sub linear space algorithm that computes approximate sliding window sums in $O(1)$ time, where the window size is given at the query time; additionally, it requires only $(1 + o(1))$ more space than is needed for a fixed window size.

1 Introduction

1.1 Background

Static dictionaries are data structures that encode a set $S \subseteq \{1, 2, \dots, n\}$ and efficiently answer *membership queries* of the form “is $i \in S$?” (for some $i \in \{1, 2, \dots, n\}$). This problem was extensively studied and memory efficient data structures that allow $O(1)$ time queries for it were suggested for several different models [6, 12, 20].

An extension of the dictionary problem is the *Rank* query, which given an identifier $i \leq n$ returns the number of elements in S that are smaller than or equal to i . For this problem as well, multiple papers proposed space efficient solutions with a constant query time [16, 21, 22]. The inverse problem, called *Select query*, asks for the ID of the i^{th} smallest element in S and was also shown to have space efficient data structures that support constant time queries [7, 21].

A seemingly different research area is the design of streaming algorithms. For many domains, such as networking, economics and databases, the ability to process large data streams is vital. As data varies over time, recent data is often considered more relevant; this motivated the study of *sliding window* algorithms, in which only the last n elements are of interest. The sliding window model was studied for many problems such as summing [3, 10, 14]; counting the number of distinct elements [2, 13]; finding frequent elements [4, 15]; answering set membership queries [5, 17, 19]; and other problems [1, 9, 18, 23]. All these works share a common goal – they significantly reduce the memory consumption; in return, they settle for approximate, rather than exact, solutions. Given sufficient space, we can solve such problems exactly simply by adding the newly arriving item into our summary and deleting the element that has left the window. However, in many applications the window size is too large and the memory requirement becomes a major bottleneck. In this paper, we show how rank queries can be used for streaming.

Even if modern RAM memories seem to be enough for storing large element sequences, there are many advantages in minimizing the memory requirements. Routers, for example, often rely on the scarce SRAM which allows access at the speed in which they are required to route packets. If the measurement algorithms are not compact enough to fit into the small SRAM, they must access the slower DRAM that does not allow real time queries. This can be a significant limitation for applications that require timely insights about the traffic, such as load balancing or denial of

service attack identification. Similarly, when implementing in software we can gain speed if we fit our algorithm into the CPU cache and reduce DRAM access. Smaller data structures might even fit in a single cache line and can be pinned there to maximize the measurement performance.

The works mentioned above significantly reduce the space requirements compared to storing the entire window in memory. However, these algorithms assume that the window size is known in advance and their data structures only allow queries about the predetermined window size to be answered efficiently. While we can maintain a different sketch for every window size that is of interest, this may be prohibitively expensive in terms of both memory and update time. Further, the goal of these algorithms is to enable memory feasible solutions to what would otherwise require storing the exact window in memory; thus, duplicating the data structures for multiple window sizes undermines the purpose for which they were created.

1.2 Our Contributions

Our first contribution is the extension of exact succinct rankers to multi sets in which every element can appear at most ℓ times. Previous works have considered multi sets under cardinality constraint for all elements combined. Here we address the natural case where every element may appear at most ℓ times, but no cardinality constraint (smaller than $n \cdot \ell$) is known for the multi set. Our approach requires $(1 + o(1))n \log(\ell + 1)$ bits and allows $O(1)$ time rank queries.

Our next contribution are novel *approximate* set and multi set representations that allow computing rank queries with an additive error of Δ , while using less space than required for storing the multi set itself. For this problem, we prove a $\mathcal{B}_{\ell, n, \Delta} = \left\lceil \frac{n}{\lceil \Delta / \ell \rceil} \right\rceil \log(\max\{\lfloor \ell / \Delta \rfloor, 1\} + 1)$ bits lower bound and propose a succinct data structure that uses $\mathcal{B}_{\ell, n, \Delta}(1 + o(1))$ bits. To the best of our knowledge, this is the first algorithm that provides approximate rank queries in $O(1)$ time using less memory than the set / multi set encoding requires.

Next, we extend the notion of approximate rankers to streams and propose algorithms that process a stream of integers in $\{0, 1, \dots, \ell\}$ and answer sliding window sum queries in $O(1)$ time. Unlike previous works [2, 3, 10], we get the window size *at query time*. That is, our algorithm can compute the sum of *any* window size while previous works assume that the size is fixed. Interestingly, our construction is succinct even when compared with the lower bound derived in [3] for fixed size windows. Thus, with a $(1 + o(1))$ space overhead we allow the algorithm to support all window sizes. This is a major improvement over the naive approach of maintaining a separate algorithm instance for every window size that is of interest, in both space and time complexity.

We note that our approach also allows approximating the sum of *historical intervals* that can be used for drill-down queries. For example, assume that we are monitoring a 100Gbps link on a backbone router such that at each second we get the utilized bandwidth (i.e., we can set $\ell = 100 \cdot 2^{30}$ bits). Now, assume that we identify a distributed denial of service attack and want to study the link utilization pattern before and during the attack. Our algorithm allows us to estimate the bandwidth between any time interval $t_1 - t_2$ (for $t_2 \leq t_1 \leq n$) simply by subtracting the estimate for the sum of the last t_2 seconds from the estimate of the last t_1 seconds' sum.

2 Related Work

2.1 Dictionaries

Consider a set $S \subseteq \{1, 2, \dots, n\}$. A *dictionary* is a data structure that supports *membership queries* of the form “Is x in S ?”. Several hashing-based works proposed methods for efficiently encoding S while supporting constant time membership queries [6, 12, 20, 24]. Dictionaries were then naturally extended to the *Indexable Dictionary* problem that also supports the operations:

1. **Rank**(i): given $i \in \{1, 2, \dots, n\}$, return $|\{y \in S : y \leq i\}|$.
2. **Select**(i): given $i \in \{1, 2, \dots, |S|\}$, return the i^{th} smallest element in S .

The problem of storing sets (and multisets, with the appropriate generalizations of the Rank and Select procedures) drew lots of attention from the research community [16, 21, 22]. Of special interest to us is the work of Jacobson [16] that allows constant time rank queries using $n + o(n)$ memory. Jacobson's idea was to look at the characteristic vector of the set, i.e., a $\{0, 1\}^n$ bits vector whose i^{th} entry is set if $i \in S$. Thus, the Rank query reduces to counting the number of set bits that precede some index i given at query time. To achieve this, Jacobson breaks the vector into $(\log n)^2$ sized *chunks*. At the end of each chunk, Jacobson keeps the number of set bits that precede it. Since there are $n/(\log n)^2$ such chunks, and each is encoded using $\log n$ bits, this requires $n/\log n = o(n)$ bits. Next, Jacobson focuses on each specific chunk and divides it into a sequence of $(1/2 \cdot \log n)$ -sized *sub-chunks*. At the end of each sub-chunk, Jacobson stores its number of preceding set bits *within the current chunk* using $O(\log \log n)$ bits. Once again, the number of sub-chunks is $n/(1/2 \cdot \log n)$ so the total memory required is $O(n \frac{\log \log n}{\log n}) = o(n)$ bits. Finally, Jacobson counts the number of set bits within each sub-chunk using a lookup table. In the table, the keys are all binary vectors of size at most $1/2 \cdot \log n$ and the values are the number of set bits; thus, the table's overall memory consumption is $O(\sqrt{n} \log n \log \log n) = o(n)$.

In this paper, we present a succinct structure for rank queries of multi sets in which each element appears at most ℓ times. This is different than the multi set representations of [20, 21] that considered cardinality constraint for the entire multi set, but without any further restriction on the number of appearances of a single item. We also provide an encoding that supports additive approximations of rank queries in less memory than required for encoding the multi set.

2.2 Algorithms that Sum over Sliding Windows

Approximating the sum of the last n elements over an integer stream, known as BASIC-SUMMING, was first introduced by Datar et al. [10]. They assumed that each element is in $\{0, 1, \dots, \ell\}$ and proposed a $(1 + \epsilon)$ multiplicative approximation algorithm. Their data structure, named Exponential Histogram (*EH*), is based on keeping timestamps of element sequences called buckets such that the last n elements fit into $O(\epsilon^{-1} \log(\ell \cdot n))$ buckets. Each bucket requires $O(\log n)$ bits to store the timestamp in addition to $O(\log \log(\ell \cdot n))$ bits to store the bucket size. Overall, the number of bits required by their algorithm is $O(\epsilon^{-1} (\log^2 n + \log \ell \cdot (\log n + \log \log \ell)))$ and it operates in amortized time $O(\frac{\log \ell}{\log n})$ or $O(\log(\ell \cdot n))$ worst case. The EH approach was then extended in [1] for other statistics over sliding windows, such as median and variance. In [14], Gibbons and Tirthapura presented a $(1 + \epsilon)$ multiplicative algorithm that operates in constant worst case time while using similar space for $\ell = n^{O(1)}$. In [3], we studied the potential memory savings one can get by replacing the $(1 + \epsilon)$ multiplicative guarantee with a Δ additive approximation. We showed that $\Theta(\frac{\ell \cdot n}{\Delta} + \log n)$ bits are required and sufficient.

In a sliding window, the last n elements get similar weight while older items do not affect the sum. Cohen and Strauss [8] considered more general aging models where older data has lower weight, but the rate in which the weight decreases may be different than that of sliding windows.

Recently, we studied [2] the affect that allowing an error in the *window size* has on the required memory of approximate summing algorithms. Specifically, we showed that if upon a query the algorithm is required to return a tuple $\langle w, \widehat{S}_w \rangle$ such that $w \in \{n, n+1, \dots, n(1+\tau)\}$ and $|\widehat{S}_w - S_w| < \Delta$ then $\Theta(\tau^{-1} \log(\frac{\tau \cdot \ell \cdot n}{\Delta}) + \log n)$ bits are needed.

All of the algorithms above assume that the window size is fixed. Here, we propose solutions that are *succinct*, even when compared to a lower bound derived here for static data, or to the bound for a fixed size window as in [3].

It is worth mentioning that these data structures *do* allow computing the sum of a window whose size is given at the query time. Alas, the query time will be slower as they do not keep aggregates that allow quick computation. Specifically, we can compute a $(1 + \epsilon)$ multiplicative approximation using a slightly extended version of EH [11] in $O(\log \epsilon^{-1} + \log \log n)$ time by a

binary search for the block with the right timestamp. We can also use the data structure of [3] for an additive approximation of Δ in $O(\min\{\frac{\ell \cdot n}{\Delta}, n\})$ time, and utilize [2]’s structure for a (τ, Δ) -approximation in time $O(\tau^{-1})$. In this paper we offer solutions that operate in $O(1)$ time.

3 Preliminaries

We say that an algorithm is *succinct* if it uses $\mathcal{B}(1 + o(1))$ bits, where \mathcal{B} is the information-theoretic lower bound for the problem it solves. Throughout the paper, we assume the standard word RAM model with a word size of $\Theta(\log n + \log \ell)$. For simplicity of presentation, we also assume that $n/(\log n)^2$ and $\sqrt{\log n}$ are integers.

► **Definition 1 (Approximation).** Given a value V and a constant $\epsilon > 0$, we say that \hat{V} is an ϵ -additive approximation of V if $V - \epsilon < \hat{V} \leq V$.¹

Next, we define the notion of an (ℓ, n, Δ) -Ranker – a structure that can answer approximate rank queries in a memory efficient manner. Specifically, $(\ell, n, 1)$ -Ranker is a succinct encoding of a multi-set over $\{0, 1, \dots, n\}$, such that no element appears more than ℓ times, that supports $O(1)$ time rank queries.

► **Definition 2 (Static Ranker).** An (ℓ, n, Δ) -Ranker, for some $\ell, n, \Delta \in \mathbb{N}^+$, is an algorithm that preprocesses a sequence in $\{0, 1, \dots, \ell\}^n$ and when queried with some $i \leq n$ returns a Δ -additive approximation \hat{S}_i of the sum of the *first* i elements, S_i , in $O(1)$ time.

We proceed with the definition of a Sliding Ranker, extending (ℓ, n, Δ) -Rankers to streams, while focusing on the *last* elements in the stream for supporting sliding window queries.

► **Definition 3 (Sliding Ranker).** An (ℓ, n, Δ) -Sliding Ranker, for some $\ell, n, \Delta \in \mathbb{N}^+$, is an algorithm that processes a stream of integers in $\{0, 1, \dots, \ell\}$ and when queried for some $i \leq n$ returns a Δ -additive approximation \hat{S}_i of the *last* i elements sum, S_i , in $O(1)$ time.

4 (ℓ, n, Δ) -Rankers

In order to construct an (ℓ, n, Δ) -Ranker, we first discuss the special case of zero-error ($\Delta = 1$).

4.1 An $(\ell, n, 1)$ -Ranker

Here, we provide a succinct construction of $(\ell, n, 1)$ -Ranker, for any ℓ, n . Intuitively, this generalizes Jacobson’s ranker [16] that addresses binary sequences ($\ell = 1$). As we show, his lookup table approach works for “small” values of ℓ . In other cases, such as $\ell \geq n$, we can split the vector into smaller and smaller intervals (i.e., sub-sub-chunk, etc.), but if the number of levels is constant, storing a lookup table for the smallest level is infeasible in $o(\mathcal{B})$ space. Thus, we use a different trick for large ℓ values for computing within-sub-chunk sums in $O(1)$. We avoid keeping the *and* the characteristic vector; instead, we keep a n -sized array in which each entry contains the sum of the sub-chunk up to that point. For example, if the sub-chunk was $\langle 1, 0, 1 \rangle$, we store $\langle 1, 1, 2 \rangle$ regardless of vector entries outside this sub-chunk. Since ℓ is “large”, this takes $o(\mathcal{B})$ space.

We start by noting that since the number of sequences in $\{0, 1, \dots, \ell\}^n$ is $(\ell + 1)^n$, any algorithm that computes such rank queries (exactly) requires $\mathcal{B} \triangleq n \log(\ell + 1)$ bits. We also note that without a query time constraint this is achievable, as we can simply store the entire data and when queried sum the required interval in $O(n)$ time. Thus, if $n = O(1)$ (i.e., we have a small array of potentially large numbers), then the same idea works and we therefore require \mathcal{B} bits; hence, we hereafter assume that $n = \omega(1)$. Next, we will prove the following:

¹ We use one-sided error, and strict inequality as this simplifies our computations.

► **Theorem 4.** *For any $\ell, n \in \mathbb{N}^+$, there exists an $(\ell, n, 1)$ -Ranker that uses $\mathcal{B}(1 + o(1))$ bits.*

We start by breaking the sequence into chunks of size $\log^2 n$, keeping the cumulative sums at the end of each chunk. The required number of bits for these sums is at most

$$\frac{n}{\log^2 n} \log(\ell \cdot n + 1) \leq \frac{n}{\log^2 n} \log((\ell + 1) \cdot n) = n \log(\ell + 1) \cdot \left(\frac{1}{\log^2 n} + \frac{1}{\log \ell \log n} \right) = o(\mathcal{B}),$$

where the last equation follows from $n = \omega(1)$.

Next, we break the chunks into sub-chunks of size $\sqrt{\log n}$ and keep the cumulative sum *from the beginning of the most recent chunk* at the each sub-chunk's end. The memory consumption of these sub-chunk aggregates is then no more than

$$\frac{n}{\sqrt{\log n}} \log(\ell \cdot \log(\log^2 n) + 1) \leq n \log(\ell + 1) \cdot \left(\frac{1}{\sqrt{\log n}} + \frac{2 \log \log n}{\log \ell \sqrt{\log n}} \right) = o(\mathcal{B}).$$

We are left with the task of efficiently computing the sub-chunk sums. Here, we split our construction depending on the relation between ℓ and n .

■ $\ell + 1 \leq 2^{\sqrt[3]{\log n}}$.

In this case, we adopt Jacobson's lookup table approach. Specifically, we create a lookup table $T : \{0, 1, \dots, \ell\}^{\sqrt{\log n}} \times \{0, 1, \dots, \sqrt{\log n} - 1\} \rightarrow \{0, 1, \dots, \ell \cdot \sqrt{\log n}\}$; the key of each table entry is a $\sqrt{\log n}$ -sized sequence of elements in $\{0, 1, \dots, \ell\}$ and an index $k \in \{0, 1, \dots, \sqrt{\log n} - 1\}$. Its value is the sum of the first k sequence entries. In order to use the table, we also store the characteristic vector itself using $n \log(\ell + 1)$ bits. The size of the table is then

$$(\ell + 1)^{\sqrt{\log n}} \sqrt{\log n} \log(\ell \sqrt{\log n} + 1) \leq 2^{\log^{5/6} n + \log \log n} \log((\ell + 1) \sqrt{\log n}) = o(\mathcal{B}).$$

Thus our overall memory consumption is $n \log(\ell + 1) \cdot (1 + o(1))$. Unfortunately, while we can consider smaller and smaller sequence aggregates, constructing such a lookup table will prevent the algorithm from being succinct when ℓ is large (e.g., for $\ell \geq n$).

■ $\ell + 1 > 2^{\sqrt[3]{\log n}}$.

In this case, we return to the cumulative approach. *Instead* of storing the characteristic vector (and without a lookup table) we store for each element the cumulative sum *from the beginning of its sub-chunk*. Since the sub-chunks are of size $\sqrt{\log n}$, the number of bits this takes is

$$n \log(\ell \sqrt{\log n} + 1) \leq n \log(\ell + 1) \cdot \left(1 + \frac{\log \log n}{\log(\ell + 1)} \right) \leq n \log(\ell + 1) \cdot \left(1 + \frac{\log \log n}{\sqrt[3]{\log n}} \right) = \mathcal{B} + o(\mathcal{B}).$$

We conclude that in all cases our construction requires $\mathcal{B}(1 + o(1))$ bits and is thus succinct.

4.2 An (ℓ, n, Δ) -Ranker for $\Delta > 1$

We start by proving a lower bound on the memory required by any (ℓ, n, Δ) -Ranker. For convenience, we denote $\mu \triangleq \Delta/\ell$. We only consider $\Delta \in \{2, \dots, \ell \cdot n\}$, as $\Delta = 1$ means zero-error and $\Delta > n \cdot \ell$ allows the algorithm to always return 0, regardless of the input.

► **Theorem 5.** *Let $\ell, n, \Delta \in \mathbb{N}^+$, then the number of bits required by any deterministic (ℓ, n, Δ) -Ranker is at least*

$$\mathcal{B}_{\ell, n, \Delta} \triangleq \lfloor n / \lceil \mu \rceil \rfloor \log(\max\{\lfloor \mu^{-1} \rfloor, 1\} + 1) = \left\lfloor \frac{n}{\lceil \Delta / \ell \rceil} \right\rfloor \log(\max\{\lfloor \ell / \Delta \rfloor, 1\} + 1).$$

Proof. We denote $I \triangleq \{\min\{\Delta \cdot k, \ell\} \mid k \in \{0, 1, \dots, \max\{\lfloor \mu^{-1} \rfloor, 1\}\}\} \subseteq \{0, 1, \dots, \ell\}$ and $\bar{I} \triangleq \{\sigma^{\lceil \mu \rceil} \mid \sigma \in I\}$. Next, consider all inputs that contain a sequence of $\lfloor n/\lceil \mu \rceil \rfloor$ blocks padded by zeros, such that each block is a member of \bar{I} ; that is, consider $\mathcal{I} \triangleq \bar{I}^{\lfloor n/\lceil \mu \rceil \rfloor} \cdot 0^{n \bmod \lceil \mu \rceil}$. Notice that each literal is in the range $\{0, 1, \dots, \ell\}$ and that each input is of size n as required. We show that every two inputs in \mathcal{I} must lead to distinct configurations in the (ℓ, n, Δ) -Ranker, thereby implying a $\lceil \log |\mathcal{I}| \rceil$ bits lower bound as required. Let $x_1 = x_{1,1}x_{1,2} \cdots x_{1,\lfloor n/\lceil \mu \rceil \rfloor} 0^{n \bmod \lceil \mu \rceil}$, $x_2 = x_{2,1}x_{2,2} \cdots x_{2,\lfloor n/\lceil \mu \rceil \rfloor} 0^{n \bmod \lceil \mu \rceil}$ be two distinct inputs in \mathcal{I} such that $x_{\alpha,\beta} \in \bar{I}$ for any $\alpha \in \{1, 2\}, \beta \in \{1, \dots, \lfloor n/\lceil \mu \rceil \rfloor\}$. Denote by $t \triangleq \min\{\gamma \in \{1, \dots, \lfloor n/\lceil \mu \rceil \rfloor\} \mid x_{1,\gamma} \neq x_{2,\gamma}\}$ the first block's index in which x_1 differs from x_2 . Now consider a query for $i \triangleq \lceil \mu \rceil \cdot t$. If $\mu \leq 1$, then $\lfloor n/\lceil \mu \rceil \rfloor = n$ and (due to the definition of I) $|x_{1,t} - x_{2,t}| \geq \Delta$, which implies an error of at least Δ for at least one of the inputs. On the other hand, $\mu > 1$ means that $I = \{0, \ell\}$ and thus either $x_{1,t} = 0^{\lceil \mu \rceil}, x_{2,t} = \ell^{\lceil \mu \rceil}$ or $x_{1,t} = \ell^{\lceil \mu \rceil}, x_{2,t} = 0^{\lceil \mu \rceil}$. In either case, the difference in sums is at least $\lceil \mu \rceil \cdot \ell \geq \Delta$. We established that if two inputs in \mathcal{I} lead to the same configuration, the error for one of them would be at least Δ while we assumed it is strictly lower. ■

We now present a succinct construction of an (ℓ, n, Δ) -Ranker. Denote $\nu \triangleq \max\{\lfloor \mu \rfloor, 1\}$, $s \triangleq \lfloor n/\nu \rfloor$ and $z \triangleq \lfloor \mu^{-1} \nu \rfloor$. For creating an (ℓ, n, Δ) -Ranker, we first show how to “compress” the input into a smaller problem that we solve exactly. Intuitively, we create a new s -long input $\bar{\rho}$, such that each of its elements is bounded by z , and then employ a $(z, s, 1)$ -Ranker, \mathcal{R} . Alas, if $\mu = \omega(1)$, this is not enough to allow succinct encoding; for this, we also compute the fraction of the input's sum that is not accounted for in $\bar{\rho}$ and use it for answering queries. Given an input $\bar{x} \in \{0, 1, \dots, \ell\}^n$, we create $\bar{\rho}$ iteratively as follows²:

$$\forall k \in \{1, 2, \dots, s\} : \quad \rho_k \triangleq \left\lfloor \Delta^{-1} \cdot \sum_{d=n-\nu \cdot k+1}^n x_d \right\rfloor - \sum_{j=1}^{k-1} \rho_j.$$

Then, we compute the remainder:

$$\mathfrak{r} \triangleq \sum_{d=1}^n x_d - \Delta \cdot \sum_{j=1}^s \rho_j. \quad (1)$$

After computing $\bar{\rho} \in \{0, 1, \dots, z\}^s$, we feed it into a $(z, s, 1)$ -Ranker denoted \mathcal{R} . Given a query for some $i \leq n$, we return²

$$\text{Query}(i) \triangleq \mathfrak{r} - (\Delta - 1/2) + \Delta \cdot \left(\sum_{j=\lfloor \frac{n-i}{\nu} \rfloor + 1}^s \rho_j \right) - \ell \cdot \rho_{\lfloor \frac{n-i}{\nu} \rfloor} \cdot (n-i) \bmod \nu, \quad$$

which we can compute in $O(1)$ as follows:³

$$\begin{aligned} \text{Query}(i) \triangleq & \mathfrak{r} - (\Delta - 1/2) + \Delta \cdot \left(\mathcal{R}.\text{Query}(s) - \mathcal{R}.\text{Query}\left(\left\lfloor \frac{n-i}{\nu} \right\rfloor\right) \right) \\ & - \ell \cdot ((n-i) \bmod \nu) \cdot \left(\mathcal{R}.\text{Query}\left(\left\lceil \frac{n-i}{\nu} \right\rceil\right) - \mathcal{R}.\text{Query}\left(\left\lfloor \frac{n-i}{\nu} \right\rfloor\right) \right). \quad (2) \end{aligned}$$

► **Lemma 6.**

$$\left(\sum_{d=1}^i x_d \right) - \Delta < \text{Query}(i) < \sum_{d=1}^i x_d.$$

² If $(n \bmod \nu) \neq 0$, we implicitly define $\rho_{\lfloor \frac{n}{\nu} \rfloor} \triangleq 0$ and $\mathcal{R}.\text{Query}\left(\left\lceil \frac{n}{\nu} \right\rceil\right) \triangleq \mathcal{R}.\text{Query}\left(\left\lfloor \frac{n}{\nu} \right\rfloor\right)$.

³ We note that if our ranker \mathcal{R} was originally constructed to compute the sum of the *last* i elements rather than the first, only two queries were needed.

Proof. We denote the error in the representation of the *last* $\nu \lfloor \frac{n-i}{\nu} \rfloor$ items by

$$\xi \triangleq \sum_{d=n-\nu \lfloor \frac{n-i}{\nu} \rfloor + 1}^n x_d - \Delta \cdot \sum_{j=1}^{\lfloor \frac{n-i}{\nu} \rfloor} \rho_j = \sum_{d=n-\nu \lfloor \frac{n-i}{\nu} \rfloor + 1}^n x_d - \Delta \left[\Delta^{-1} \cdot \sum_{d=n-\nu \lfloor \frac{n-i}{\nu} \rfloor + 1}^n x_d \right].$$

Observe that $\xi = \left(\sum_{d=n-\nu \lfloor \frac{n-i}{\nu} \rfloor + 1}^n x_d \bmod \Delta \right)$ and hence

$$0 \leq \xi \leq \Delta - 1. \quad (3)$$

Next, we use (1) to obtain

$$\begin{aligned} \sum_{d=1}^i x_d &= \sum_{d=1}^n x_d - \sum_{d=i+1}^n x_d = \mathbf{r} + \Delta \cdot \sum_{j=1}^s \rho_j - \sum_{d=i+1}^n x_d \\ &= \mathbf{r} + \Delta \cdot \left(\sum_{j=1}^{\lfloor \frac{n-i}{\nu} \rfloor} \rho_j + \sum_{j=\lfloor \frac{n-i}{\nu} \rfloor + 1}^s \rho_j \right) - \sum_{d=i+1}^{n-\nu \lfloor \frac{n-i}{\nu} \rfloor} x_d - \sum_{d=n-\nu \lfloor \frac{n-i}{\nu} \rfloor + 1}^n x_d \\ &= \mathbf{r} - \xi + \Delta \cdot \sum_{j=\lfloor \frac{n-i}{\nu} \rfloor + 1}^s \rho_j - \sum_{d=i+1}^{n-\nu \lfloor \frac{n-i}{\nu} \rfloor} x_d \\ &= \text{Query}(i) - \xi - \sum_{d=i+1}^{n-\nu \lfloor \frac{n-i}{\nu} \rfloor} x_d + \Delta - 1/2 + \ell \cdot \rho_{\lfloor \frac{n-i}{\nu} \rfloor} \cdot (n-i) \bmod \nu. \end{aligned} \quad (4)$$

We now perform a case analysis, based on the value of μ and start with the simpler case where $\mu < 2$. In this case, we have $\nu = 1$ and thus we can rearrange (4) as:

$$\text{Query}(i) - \sum_{d=1}^i x_d = \xi - \Delta + 1/2,$$

and using (3) we immediately get $\left(\sum_{d=1}^i x_d \right) - \Delta < \text{Query}(i) < \sum_{d=1}^i x_d$.

Next, we focus on the case of $\mu \geq 2$. Thus, we hereafter have $\nu = \lfloor \mu \rfloor$ and $\forall j \in \{1, 2, \dots, s\} : \bar{\rho}_j \in \{0, 1\}$. We now consider if and when both $\rho_{\lfloor \frac{n-i}{\nu} \rfloor} = 1$ and $(n-i) \bmod \nu \neq 0$ (which implies $\lfloor \frac{n-i}{\nu} \rfloor = \lfloor \frac{n-i}{\nu} \rfloor + 1$); observe that

$$\begin{aligned} \rho_{\lfloor \frac{n-i}{\nu} \rfloor} &= \left[\Delta^{-1} \cdot \sum_{d=n-\nu \cdot \lfloor \frac{n-i}{\nu} \rfloor + 1}^n x_d \right] - \sum_{j=1}^{\lfloor \frac{n-i}{\nu} \rfloor - 1} \rho_j = \left[\Delta^{-1} \cdot \sum_{d=n-\nu \cdot \lfloor \frac{n-i}{\nu} \rfloor + 1}^n x_d - \sum_{j=1}^{\lfloor \frac{n-i}{\nu} \rfloor} \rho_j \right] \\ &= \left[\Delta^{-1} \cdot \sum_{d=n-\nu \cdot \lfloor \frac{n-i}{\nu} \rfloor + 1}^n x_d + \Delta^{-1} \cdot \left(\xi - \sum_{d=n-\nu \cdot \lfloor \frac{n-i}{\nu} \rfloor + 1}^n x_d \right) \right] = \left[\Delta^{-1} \cdot \left(\xi + \sum_{d=n-\nu \cdot \lfloor \frac{n-i}{\nu} \rfloor + 1}^{n-\nu \cdot \lfloor \frac{n-i}{\nu} \rfloor} x_d \right) \right]. \end{aligned}$$

Thus, if $(n-i) \bmod \nu \neq 0$, then

$$\rho_{\lfloor \frac{n-i}{\nu} \rfloor} = 1 \iff \xi + \sum_{d=n-\nu \lfloor \frac{n-i}{\nu} \rfloor + 1}^{n-\nu \cdot \lfloor \frac{n-i}{\nu} \rfloor} x_d \geq \Delta \iff \xi + \sum_{d=i+1}^{n-\nu \cdot \lfloor \frac{n-i}{\nu} \rfloor} x_d \geq \Delta - \sum_{d=n-\nu \lfloor \frac{n-i}{\nu} \rfloor + 1}^i x_d. \quad (5)$$

Next, we split to cases based on the value of $\rho_{\lfloor \frac{n-i}{\nu} \rfloor}$:

■ $\rho_{\lceil \frac{n-i}{\nu} \rceil} = 1$. In this case, according to (4) and (5) we have:

$$\begin{aligned}
 \text{Query}(i) - \sum_{d=1}^i x_d &= \xi + \sum_{d=i+1}^{n-\nu \lfloor \frac{n-i}{\nu} \rfloor} x_d - (\Delta - 1/2 + \ell \cdot (n-i) \bmod \nu) \\
 &\geq \Delta - \sum_{d=n-\nu \lfloor \frac{n-i}{\nu} \rfloor + 1}^i x_d - (\Delta - 1/2 + \ell \cdot (n-i) \bmod \nu) \\
 &= - \sum_{d=n-\nu \lfloor \frac{n-i}{\nu} \rfloor + 1}^i x_d + 1/2 - \ell \cdot (n-i) \bmod \nu \\
 &\geq - \sum_{d=n-\nu \lfloor \frac{n-i}{\nu} \rfloor + 1}^i \ell + 1/2 - \ell \cdot (n-i) \bmod \nu \\
 &\geq -(\ell\nu - 1) + 1/2 = -(\ell \lfloor \Delta/\ell \rfloor - 1) + 1/2 \geq -\Delta + 1/2.
 \end{aligned}$$

On the other hand, we bound the error from above as follows:

$$\begin{aligned}
 \text{Query}(i) - \sum_{d=1}^i x_d &= \xi + \sum_{d=i+1}^{n-\nu \lfloor \frac{n-i}{\nu} \rfloor} x_d - (\Delta - 1/2 + \ell \cdot (n-i) \bmod \nu) \\
 &\leq \Delta - 1 + \sum_{d=i+1}^{n-\nu \lfloor \frac{n-i}{\nu} \rfloor} \ell - \Delta - \ell \cdot (n-i) \bmod \nu + 1/2 \leq -1/2.
 \end{aligned}$$

■ $\rho_{\lceil \frac{n-i}{\nu} \rceil} = 0$. Similarly to before, using (4) and (5) we get:

$$\begin{aligned}
 \text{Query}(i) - \sum_{d=1}^i x_d &= \xi + \sum_{d=i+1}^{n-\nu \lfloor \frac{n-i}{\nu} \rfloor} x_d - (\Delta - 1/2) < \Delta - \sum_{d=n-\nu \lfloor \frac{n-i}{\nu} \rfloor + 1}^i x_d - (\Delta - 1/2) \\
 &= - \sum_{d=n-\nu \lfloor \frac{n-i}{\nu} \rfloor + 1}^i x_d + 1/2 \leq 1/2.
 \end{aligned}$$

Now, we use the fact that both $\text{Query}(i)$ and $\sum_{d=1}^i x_d$ are integers to deduce that $\text{Query}(i) - \sum_{d=1}^i x_d < 1/2 \implies \text{Query}(i) - \sum_{d=1}^i x_d \leq 0$. Finally, we bound the error from above:

$$\text{Query}(i) - \sum_{d=1}^i x_d = \xi + \sum_{d=i+1}^{n-\nu \lfloor \frac{n-i}{\nu} \rfloor} x_d - (\Delta - 1/2) \geq -\Delta + 1/2.$$

We conclude that in all cases we have $\left(\sum_{d=1}^i x_d\right) - \Delta < \text{Query}(i) < \sum_{d=1}^i x_d$. ■

Next, we show that the value of each entry in $\bar{\rho}$ is smaller than z , as stated.

► **Lemma 7.** For any $k \in \{1, 2, \dots, s\}$, $\rho_k \leq \lfloor \mu^{-1} \nu \rfloor$.

Proof. Notice that $\rho_1 = \lfloor \Delta^{-1} \cdot \sum_{d=1}^{\nu} x_d \rfloor \leq \lfloor \Delta^{-1} \cdot \nu \ell \rfloor = \lfloor \mu^{-1} \nu \rfloor$. For other k values, we have

$$\begin{aligned}
 \rho_k &= \left\lfloor \Delta^{-1} \cdot \sum_{d=1}^{\nu \cdot k} x_d \right\rfloor - \sum_{j=1}^{k-1} \rho_j = \left\lfloor \Delta^{-1} \cdot \sum_{d=\nu \cdot (k-1)+1}^{\nu \cdot k} x_d + \left(\Delta^{-1} \cdot \sum_{d=1}^{\nu \cdot (k-1)} x_d - \sum_{j=1}^{k-2} \rho_j \right) \right\rfloor - \rho_{k-1} \\
 &\leq \left\lfloor \Delta^{-1} \cdot \sum_{d=\nu \cdot (k-1)+1}^{\nu \cdot k} x_d + \left(\left\lfloor \Delta^{-1} \cdot \sum_{d=1}^{\nu \cdot (k-1)} x_d \right\rfloor - \sum_{j=1}^{k-2} \rho_j \right) \right\rfloor - \rho_{k-1}
 \end{aligned}$$

$$= \left\lfloor \Delta^{-1} \cdot \sum_{d=\nu \cdot (k-1)+1}^{\nu \cdot k} x_d \right\rfloor \leq \lfloor \Delta^{-1} \cdot \nu \ell \rfloor = \lfloor \mu^{-1} \nu \rfloor. \quad \blacksquare$$

We now bound \mathfrak{r} for analyzing the space of our construction; the proof appears in Appendix A.

► **Lemma 8.** *For any input $x \in \{0, 1, \dots, \ell\}^n$, the remainder in (1) satisfies $\mathfrak{r} < 2\Delta$.*

Follows is an analysis of our ranker.

► **Lemma 9.** *Let $\ell, n, \Delta \in \mathbb{N}^+$ and $\mu \triangleq \Delta/\ell$. The number of bits required by our ranker is $(1 + o(1)) \cdot \lfloor n / \max\{\lfloor \mu \rfloor, 1\} \rfloor \cdot \log(\lceil \mu^{-1} \rceil + 1)$.*

Proof. Our construction has two components: the exact ranker \mathcal{R} and the remainder \mathfrak{r} . As \mathcal{R} is a $(z, s, 1)$ -Ranker, where $s \triangleq \lfloor n/\nu \rfloor$ and $z \triangleq \lfloor \mu^{-1} \nu \rfloor$, it requires $(1 + o(1)) \cdot s \log(z + 1)$ bits according to Theorem 4. Recalling that $\nu = \max\{\lfloor \mu \rfloor, 1\}$ gives us the desired $\mathcal{B}_{\ell, n, \Delta}(1 + o(1))$ bound. Finally, Lemma 8 tells us that $\mathfrak{r} < 2\Delta$ and can therefore be represented using $O(\log \Delta) = o(\mathcal{B}_{\ell, n, \Delta})$ bits. \blacksquare

► **Theorem 10.** *Let $\ell, n, \Delta \in \mathbb{N}^+$ such that $(\mu = o(1)) \vee (\mu = \omega(1)) \vee (\mu \in \mathbb{N}) \vee (\mu^{-1} \in \mathbb{N})$, the construction above is an (ℓ, n, Δ) -Ranker that uses $\mathcal{B}_{\ell, n, \Delta}(1 + o(1))$ bits.⁴*

Proof. Recall that $\mathcal{B}_{\ell, n, \Delta} = \lfloor n / \lceil \mu \rceil \rfloor \log(\max\{\lfloor \mu^{-1} \rfloor, 1\} + 1)$ while our algorithm uses $(1 + o(1)) \cdot \lfloor n / \max\{\lfloor \mu \rfloor, 1\} \rfloor \cdot \log(\lceil \mu^{-1} \rceil + 1)$ bits. If $\mu = o(1)$, we have $\mathcal{B}_{\ell, n, \Delta} = n \log(\lfloor \mu^{-1} \rfloor + 1) = (1 - o(1))n \log \mu^{-1}$ when our structure takes $(1 + o(1)) \cdot n \cdot \log(\lceil \mu^{-1} \rceil + 1) = (1 + o(1))n \log \mu^{-1}$ bits. Similarly, if $\mu = \omega(1)$ then $\mathcal{B}_{\ell, n, \Delta} = \lfloor n / \lceil \mu \rceil \rfloor = (1 - o(1)) \cdot (n/\mu)$ while we require $(1 + o(1)) \cdot \lfloor n / \lfloor \mu \rfloor \rfloor = (1 + o(1)) \cdot (n/\mu)$. The case for $(\mu = \Theta(1)) \wedge ((\mu \in \mathbb{N}) \vee (\mu^{-1} \in \mathbb{N}))$ follows from similar arguments. \blacksquare

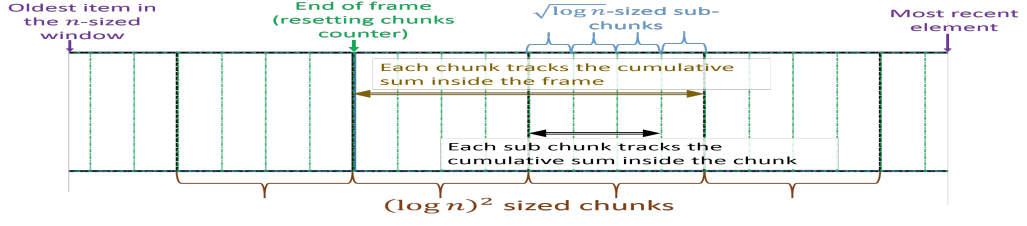
5 (ℓ, n, Δ) -Sliding Rankers

As in the case of static data rankers, we first consider the exact case where $\Delta = 1$.

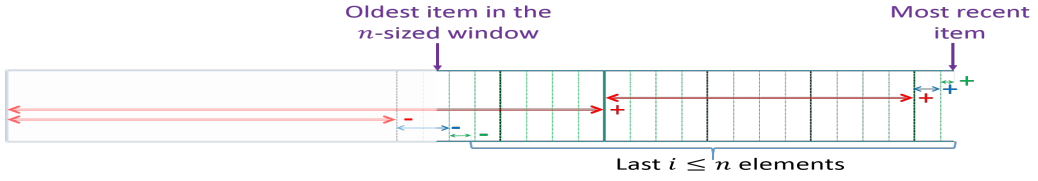
5.1 An $(\ell, n, 1)$ -Sliding Ranker

In this section, we provide a construction for an $(\ell, n, 1)$ -Sliding Ranker that requires $\mathcal{B}(1 + o(1))$ bits, where $\mathcal{B} \triangleq n \log(\ell + 1)$ is the information-theoretic lower bound even without considering sliding windows. Intuitively, we adapt our $(\ell, n, 1)$ -Ranker construction to the sliding window setting by incrementally building the chunks and sub-chunks. We start by breaking the stream into n -sized *frames*. As in the original construction, we split the frames into $(\log n)^2$ sized *chunks*, where each chunk is further divided into $\sqrt{\log n}$ -sized *sub-chunks*. In the case where $\ell + 1 \leq 2^{\sqrt[3]{\log n}}$, we keep a $O(2^{\log^{5/6} n + \log \log n} \log(\ell \sqrt{\log n}))$ sized lookup table that maps each sequence in $\{0, 1, \dots, \ell\}^{\leq \sqrt{\log n}}$ to its sum. If $\ell + 1 > 2^{\sqrt[3]{\log n}}$, we simply track the sums *within a sub-chunk* by keeping the cumulative sum for each item. We keep the chunk aggregates in a $n/(\log n)^2$ -sized circular buffer, and the sub-chunk aggregates in a similar structure of size $n/\sqrt{\log n}$. Finally, we “reset” the frame accumulator every n elements, so that each chunk’s aggregate is always smaller than $n \cdot \ell$. Since each of the chunk aggregates requires $O(\log(\ell n))$ bits and each of the sub-chunk aggregates takes $O(\log(\ell \log n))$ bits, our overall space consumption is as required. Our $(\ell, n, 1)$ -Sliding Ranker construction is illustrated in Figure 1, while the query procedure is exemplified in Figure 2. In Appendix D we provide an algorithm for the $\ell + 1 > 2^{\sqrt[3]{\log n}}$ case; here, we hereafter assume that $\ell + 1 \leq 2^{\sqrt[3]{\log n}}$. Our algorithm uses the following variables:

⁴ In other cases, our construction uses at most $B(2 + o(1))$ bits but might not be succinct.



■ **Figure 1** The $(\ell, n, 1)$ -Sliding Ranker construction of Algorithm 1. We split the stream into frames, the frames into chunks, and the chunks into sub-chunks. At the end of each chunk we keep the sum of all elements that preceded it in the frame. Similarly, for each sub-chunk we keep the sum of items from the beginning of its chunk.



■ **Figure 2** A query example of Algorithm 1. To compute the sum of the green ranges, we use either a lookup table or the within-sub-chunk aggregates depending on whether $\ell + 1 \leq 2\sqrt[3]{\log n}$. The sub-chunks aggregate array allows us to retrieve the sum of the blue intervals, and the chunk aggregates contain the sum of the red ranges. We first add values to get the sum of elements from the beginning of the frame in which the i -long interval starts. Then, we subtract the sum of items that arrived more than i elements ago.

- C - a cyclic buffer of $n/(\log n)^2$ integers, each allocated with $\lceil \log(\ell n + 1) \rceil$ bits.
- SC - a cyclic buffer of $n/\sqrt{\log n}$ integers, each allocated with $\lceil \log(\ell \log^2 n + 1) \rceil$ bits.
- $total$ - the sum of elements inside the current frame.
- ind - the index of the most recent item, modulo n .
- T - a lookup table mapping sequences of length $\leq \sqrt{\log n}$ to their sums.
- \mathcal{W} - the last n elements window.

We give a pseudo code of our $(\ell, n, 1)$ -Sliding Ranker in Algorithm 1. We now formulate the properties of the algorithm; the theorem's proof is deferred to Appendix B due to lack of space.

► **Theorem 11.** *Algorithm 1 is an $(\ell, n, 1)$ -Sliding Ranker that uses $\mathcal{B}(1 + o(1))$ memory bits.*

5.2 An (ℓ, n, Δ) -Sliding Ranker for $\Delta > 1$

Similarly to the way we used $(\ell, n, 1)$ -Rankers to construct (ℓ, n, Δ) -Rankers for any $\Delta \in \{2, \dots, \ell \cdot n\}$, we now use the exact $(\ell, n, 1)$ -Sliding Ranker for constructing an (ℓ, n, Δ) -Sliding Ranker.

Intuitively, we split the stream into *blocks* of size ν and construct the remainder \mathbf{r} gradually; whenever a block ends, we compute a new ρ_k value and feed it into an exact $(z, s, 1)$ -Sliding Ranker we use as a black box. When queried, we employ our exact ranker and remainder to estimate the relevant sum, similarly to our (ℓ, n, Δ) -Ranker queries from Section 4.2. However, if we simply sum the elements using \mathbf{r} , it will require $\Omega(\log \ell)$ bits; this will not allow us to remain succinct if $\mu = \Omega(1)$ as the lower bound for this case is $\mathcal{B}_{\ell, n, \Delta} = O(n)$ and is independent of ℓ (given that $\mu = \Delta/\ell$ is fixed). To solve this, we follow [3]'s approach and round every arriving element, representing it using $\mathbf{b} \triangleq \lceil \log(n/\mu) + \log \log n \rceil$ bits. That is, if $x \in \{0, 1, \dots, \ell\}$ arrived, we consider $\text{Round}_{\mathbf{b}}(x) \triangleq 2^{-\mathbf{b}} \ell \cdot \left\lceil \frac{x 2^{\mathbf{b}}}{\ell} \right\rceil$ instead. To compensate for the rounding error, we will need blocks of size smaller than that we used in our (ℓ, n, Δ) -Ranker construction; specifically,

Algorithm 1 An $(\ell, n, 1)$ -Sliding Ranker for $\ell + 1 \leq 2^{\sqrt[3]{\log n}}$

Initialization: $C \leftarrow \bar{0}, SC \leftarrow \bar{0}, ind \leftarrow 0, total \leftarrow 0, \mathcal{W} \leftarrow \bar{0}$
 $T \leftarrow \left(\{0, 1, \dots, \ell\}^{\leq \sqrt{\log n}} \rightarrow \{0, 1, \dots, \ell \cdot \sqrt{\log n}\} \right)$ lookup table

- 1: **function** ADD(element x)
- 2: $ind \leftarrow ind + 1 \bmod n$
- 3: $\mathcal{W}[ind] \leftarrow x$
- 4: **if** $(ind \bmod \sqrt{\log n}) = 0$ **then** ▷ End of a sub-chunk
- 5: $sum \leftarrow T \left[\mathcal{W}[(1 + ind - \sqrt{\log n}) \bmod n, \dots, ind] \right]$ ▷ The sub-chunk's sum
- 6: $total \leftarrow total + sum$
- 7: $SC[ind/\sqrt{\log n}] \leftarrow sum$
- 8: **if** $(ind \bmod (\log n)^2) = 0$ **then** ▷ End of a chunk
- 9: $C[ind/(\log n)^2] \leftarrow total$
- 10: **if** $(ind \bmod n) = 0$ **then** ▷ End of a frame, reset counter
- 11: $total \leftarrow 0$
- 12: **function** QUERY(i) ▷ For $i \leq n$
- 13: $addition \leftarrow 0$
- 14: **if** $i \geq ind$ **then** ▷ If not contained in current frame
- 15: $addition \leftarrow C[0]$
- 16: **return** ▷ See Figure 2

$$T \left[\mathcal{W} \left[ind - (ind \bmod \sqrt{\log n}) + 1, \dots, ind \right] \right]$$

$$+ SC \left[\left\lfloor ind/\sqrt{\log n} \right\rfloor \right] + C \left[\left\lfloor ind/(\log n)^2 \right\rfloor \right] + addition$$

$$- C \left[\left\lfloor ((ind - i) \bmod n) / (\log n)^2 \right\rfloor \right] - SC \left[\left\lfloor ((ind - i) \bmod n) / \sqrt{\log n} \right\rfloor \right]$$

$$- T \left[\mathcal{W} \left[\left\lfloor ((ind - i) \bmod n) / \sqrt{\log n} \right\rfloor + 1, \dots, ind - i \right] \right]$$

we set $\nu \triangleq \max \{ \lfloor \mu \cdot (1 - 1/\log n) \rfloor, 1 \}$. Additionally, when $\mu < 1$ the block size has to remain 1, so we have to compensate for the rounding error by other means; this is achieved by reducing the “sensitivity” to $\tilde{\Delta} \triangleq \lfloor \Delta \cdot (1 - 1/\log n) \rfloor$.⁵ The parameters for the exact ranker are then $s \triangleq \lfloor n/\nu \rfloor$ and $z \triangleq \lfloor \mu^{-1}\nu \rfloor$. Our algorithm uses the following variables:

- \mathfrak{R} - a $(z, s, 1)$ -Sliding Ranker, as described in Section 5.1.
- \mathfrak{r} - tracks the sum of elements that is not yet recorded in \mathfrak{R} .
- o - the offset within the block.

A pseudo code of our method appears in Algorithm 2. Next follows a memory analysis of the algorithm with a proof given in Appendix C.

► **Lemma 12.** *Algorithm 2 requires $(1 + o(1)) \cdot \lfloor n / \max \{ \lfloor \mu \rfloor, 1 \} \rfloor \cdot \log(\lceil \mu^{-1} \rceil + 1) + O(\log n)$ bits.*

This allows us to conclude, similarly to Theorem 10, that our algorithm is succinct if the error satisfies $\Delta = o\left(\frac{\ell \cdot n}{\log n}\right)$. We also note that a $\lfloor \log n \rfloor$ lower bound was shown in [3] even when only

⁵ If $\tilde{\Delta} = 1$, then we simply apply the exact algorithm from the previous subsection.

Algorithm 2 An (ℓ, n, Δ) -Sliding Ranker algorithm

```

1: Initialization:  $\mathbf{r} \leftarrow 0, o \leftarrow 0, \mathfrak{R} \leftarrow (z, s, 1) \leftarrow \text{Sliding Ranker.init}()$ 
2: function ADD(ELEMENT  $x$ )
3:    $o \leftarrow (o + 1) \bmod \nu$ 
4:    $\mathbf{r} \leftarrow \mathbf{r} + \text{Round}_b(x)$ 
5:   if  $o = 0$  then
6:      $\rho \leftarrow \left\lfloor \tilde{\Delta}^{-1} \cdot \mathbf{r} \right\rfloor$ 
7:      $\mathbf{r} \leftarrow \mathbf{r} - \tilde{\Delta} \cdot \rho$ 
8:      $\mathfrak{R}.\text{ADD}(\rho)$ 
9: function QUERY( $i$ )
10:  if  $i \leq o$  then
11:    return  $\mathbf{r} - (\tilde{\Delta} - 1/2)$ 
12:  else
13:     $\text{numElems} \leftarrow \left\lceil \frac{i-o}{\nu} \right\rceil$ 
14:     $\text{totalSum} \leftarrow \mathfrak{R}.\text{QUERY}(\text{numElems})$ 
15:     $\text{oldest}_\rho \leftarrow \text{totalSum} - \mathfrak{R}.\text{QUERY}(\text{numElems} - 1)$ 
16:     $\text{out} \leftarrow (\nu - ((i - o) \bmod \nu))$ 
17:    return  $\mathbf{r} - (\tilde{\Delta} - 1/2) + \tilde{\Delta} \cdot \text{totalSum} - \ell \cdot \text{oldest}_\rho \cdot \text{out}$ 

```

fixed sized windows (where $i \equiv n$) are considered. Thus, our algorithm always requires at most $O(\mathcal{B}_{\ell, n, \Delta})$, even if the allowed error is $\Omega\left(\frac{\ell \cdot n}{\log n}\right)$.

► **Corollary 13.** Let $\ell, n, \Delta \in \mathbb{N}^+$ such that $\mu \triangleq \Delta/\ell$ satisfies

$$\left(\mu = o\left(\frac{n}{\log n}\right)\right) \wedge [(\mu = o(1)) \vee (\mu = \omega(1)) \vee (\mu \in \mathbb{N}) \vee (\mu^{-1} \in \mathbb{N})],$$

then Algorithm 2 is succinct. For other parameters, it uses $O(\mathcal{B}_{\ell, n, \Delta})$ space.

The following theorem, whose proof is deferred to Appendix E due to lack of space, shows the correctness of the algorithm.

► **Theorem 14.** Algorithm 2 is an (ℓ, n, Δ) -Sliding Ranker.

6 Discussion

In this paper, we studied the properties of data structures that support *approximate* rank queries for multi sets in which each element in $\{1, 2, \dots, n\}$ appears at most ℓ times. We showed a lower bound for the problem and succinct constructions that require $(1 + o(1))$ times as much memory. We then extended our approach and provided algorithms that process data streams and handle sliding window sum queries. Unlike previous work, we do not assume that the window size is fixed but rather get it at the query time. Interestingly, we show that this is doable in constant time and an additional $(1 + o(1))$ space factor.

In the future, we would like to study structures that allow approximate select queries in $O(1)$ time. This will allow efficient approximate-percentile computation for multi sets. We note that this is already achievable with our data structure in $O(\log n)$ time using a binary search over the rank queries. We also plan to explore the possibility of creating approximate rankers with a *multiplicative* error rather than additive. Finally, we wish to extend our approach to problems other than summing; e.g., computing heavy hitters for a sliding window whose size is given at the query time.

References

- 1 Brian Babcock, Mayur Datar, Rajeev Motwani, and Liadan O’Callaghan. Maintaining variance and k-medians over data stream windows. In *ACM PODS*, 2003.
- 2 Ran Ben Basat, Gil Einziger, and Roy Friedman. Efficient network measurements through approximated windows. *CoRR*:1703.01166.
- 3 Ran Ben Basat, Gil Einziger, Roy Friedman, and Yaron Kassner. Efficient Summing over Sliding Windows. In *SWAT*, 2016.
- 4 Ran Ben-Basat, Gil Einziger, Roy Friedman, and Yaron Kassner. Heavy hitters in streams and sliding windows. In *IEEE INFOCOM*, 2016.
- 5 Ran Ben-Basat, Gil Einziger, Roy Friedman, and Yaron Kassner. Poster abstract: A sliding counting bloom filter. In *IEEE INFOCOM*, 2017.
- 6 Andrej Brodnik and J Ian Munro. Membership in constant time and almost-minimum space. *SIAM Journal on Computing*, 28(5):1627–1640, 1999.
- 7 David Clark. *Compact Pat trees*. PhD thesis, PhD thesis, University of Waterloo, 1998.
- 8 Edith Cohen and Martin J. Strauss. Maintaining time-decaying stream aggregates. *Journal of algorithms*, 59(1):19–36, 2006.
- 9 Michael S Crouch, Andrew McGregor, and Daniel Stubbs. Dynamic graphs in the sliding-window model. In *ESA*. 2013.
- 10 Mayur Datar, Aristides Gionis, Piotr Indyk, and Rajeev Motwani. Maintaining stream statistics over sliding windows. *SIAM Journal of Computing*.
- 11 Mayur Datar, Aristides Gionis, Piotr Indyk, and Rajeev Motwani. Maintaining stream statistics over sliding windows. *SIAM Journal of Computing*, 31(6):1794–1813, 2002.
- 12 Michael L Fredman, János Komlós, and Endre Szemerédi. Storing a sparse table with 0 (1) worst case access time. *Journal of the ACM*, 31(3):538–544, 1984.
- 13 Éric Fusy and Frédéric Giroire. Estimating the number of active flows in a data stream over a sliding window. In *ANALCO*, 2007.
- 14 Phillip B. Gibbons and Srikanta Tirthapura. Distributed streams algorithms for sliding windows. In *SPAA*, 2002.
- 15 Regant Y. S. Hung, Lap-Kei Lee, and Hing-Fung Ting. Finding frequent items over sliding windows with constant update time. *Information Proceedings Letters10’*, 110(7):257–260.
- 16 Guy Joseph Jacobson. *Succinct Static Data Structures*. PhD thesis, 1988.
- 17 Yang Liu, Wenji Chen, and Yong Guan. Near-optimal approximate membership query over time-decaying windows. In *IEEE INFOCOM*, 2013.
- 18 Kyriakos Mouratidis, Spiridon Bakiras, and Dimitris Papadias. Continuous monitoring of top-k queries over sliding windows. In *ACM SIGMOD*, 2006.
- 19 Moni Naor and Eylon Yogev. Sliding bloom filters. In *Algorithms and Computation*. 2013.
- 20 Rasmus Pagh. Low redundancy in static dictionaries with constant query time. *SIAM Journal on Computing*, 31(2):353–363, 2001.
- 21 Rajeev Raman, Venkatesh Raman, and Srinivasa Rao Satti. Succinct indexable dictionaries with applications to encoding k-ary trees, prefix sums and multisets. *ACM Trans. Algorithms*, 3(4):43, 2007.
- 22 Venkatesh Raman and S. Srinivasa Rao. Static dictionaries supporting rank. In *ISAAC*, 1999.
- 23 Zhitao Shen, M.A. Cheema, Xuemin Lin, Wenjie Zhang, and Haixun Wang. Efficiently monitoring top-k pairs over sliding windows. In *IEEE ICDE*, 2012.
- 24 Robert Endre Tarjan and Andrew Chi-Chih Yao. Storing a sparse table. *Communications of the ACM*, 22(11):606–611, 1979.

A

 Proof of Lemma 8

Proof.

$$\begin{aligned}
\tau &= \sum_{d=1}^n x_d - \Delta \cdot \sum_{j=1}^s \rho_j = \sum_{d=1}^n x_d - \Delta \cdot \left(\left\lfloor \Delta^{-1} \cdot \sum_{d=n-\nu \cdot s+1}^n x_d \right\rfloor - \sum_{j=1}^{s-1} \rho_j \right) - \Delta \cdot \sum_{j=1}^{s-1} \rho_j \\
&= \sum_{d=1}^n x_d - \Delta \cdot \left\lfloor \Delta^{-1} \cdot \sum_{d=n-\nu \cdot \lfloor n/\nu \rfloor + 1}^n x_d \right\rfloor \leq \Delta - 1 + \sum_{d=1}^{n-\nu \cdot \lfloor n/\nu \rfloor} x_d \\
&\leq \Delta - 1 + \sum_{d=1}^{\nu-1} x_d \leq \Delta - 1 + (\nu - 1)\ell.
\end{aligned}$$

If $\mu \leq 1$, then $\nu = 1$ and thus $\tau \leq \Delta - 1$. Otherwise, we have $\tau \leq \Delta - 1 + (\mu - 1)\ell \leq 2\Delta - \ell - 1$. ■

B

 Proof of Theorem 11

We start with analyzing the memory requirements of our algorithm.

► **Lemma 15.** *Algorithm 1 uses $\mathcal{B}(1 + o(1))$ memory bits.*

Proof. We have $n/(\log n)^2$ chunks, each represented using $O(\log(\ell n))$ bits. Similarly, each of the $n/\sqrt{\log n}$ sub-chunk aggregates requires $O(\log(\ell \log n))$ bits as its value is bounded by $\log(\ell \log^2 n)$. Our window, \mathcal{W} uses $n \log(\ell + 1)$ bits, while the *total* and *ind* variables require $O(\log n)$ bits. Thus, the overall space consumption is $\mathcal{B}(1 + o(1))$. ■

We are now ready to prove the theorem.

Proof. Denote the stream by $x_1, x_2, \dots, x_{k \cdot n + m}$, such that the most recent element's index is $k \cdot n + m$, where $m \in [n - 1]$ is the offset within the current frame and k frames were completed so far. We assume that $k \geq 1$. The case for $k = 0$ follows from similar arguments. We start with a few straight forward observations. Notice that $C[0]$ always contains the sum of the last frame that was completed; that is, $C[0] = \sum_{d=(k-1)n+1}^{k \cdot n} x_d$. Next, for any positive $j \in [n/(\log n)^2 - 1]$, we have that $C[j]$ contains the sum of the last j -indexed chunk that was completed, i.e.,

$$C[j] = \begin{cases} \sum_{d=k \cdot n + j \cdot (\log n)^2}^{k \cdot n + j \cdot (\log n)^2 + 1} x_d & \text{if } m \geq j \cdot (\log n)^2 \\ \sum_{d=(k-1)n + j \cdot (\log n)^2}^{(k-1)n + j \cdot (\log n)^2 + 1} x_d & \text{otherwise} \end{cases}.$$

Similarly, we have that $\forall j \in [2n/\log n]$:

$$SC[j] = \begin{cases} \sum_{d=k \cdot n + j \cdot \sqrt{\log n}}^{k \cdot n + j \cdot \sqrt{\log n} + 1} x_d & \text{if } m \geq j \cdot \sqrt{\log n} \\ \sum_{d=(k-1)n + j \cdot \sqrt{\log n}}^{(k-1)n + j \cdot \sqrt{\log n} + 1} x_d & \text{otherwise} \end{cases}.$$

Given a query for $i \leq n$, the goal of an $(\ell, n, 1)$ -Sliding Ranker is to return the quantity $S \triangleq \sum_{d=k \cdot n + m - i + 1}^{k \cdot n + m} x_d$. First, we express the sum of elements from the beginning of the *previous* frame, S_P , as:

$$S_P \triangleq \sum_{d=(k-1)n+1}^{k \cdot n + m} x_d = \sum_{d=(k-1)n+1}^{kn} x_d + \sum_{d=kn+1}^{\lfloor (k \cdot n + m)/(\log n)^2 \rfloor} x_d + \sum_{d=\lfloor (k \cdot n + m)/(\log n)^2 \rfloor + 1}^{\lfloor (k \cdot n + m)/\sqrt{\log n} \rfloor} x_d + \sum_{d=\lfloor (k \cdot n + m)/\sqrt{\log n} \rfloor + 1}^{k \cdot n + m} x_d.$$

Next, since $\text{ind} = (k \cdot n + m \bmod n) = m$, we have that

1. $C[0] = \sum_{d=(k-1)n+1}^{k \cdot n+m} x_d.$
2. $C \left[\left\lfloor \frac{ind}{(\log n)^2} \right\rfloor \right] = \sum_{d=kn+1}^{\left\lfloor \frac{ind}{(\log n)^2} \right\rfloor} x_d.$
3. $SC \left[\left\lfloor \frac{ind}{\sqrt{\log n}} \right\rfloor \right] = \sum_{d=\left\lfloor \frac{ind}{(\log n)^2} \right\rfloor+1}^{\left\lfloor \frac{(k \cdot n+m)}{\sqrt{\log n}} \right\rfloor} x_d.$
4. $T \left[x_{\left\lfloor \frac{ind}{\sqrt{\log n}} \right\rfloor+1}, \dots, x_{k \cdot n+m} \right] = \sum_{d=\left\lfloor \frac{(k \cdot n+m)}{\sqrt{\log n}} \right\rfloor+1}^{k \cdot n+m} x_d.$

Notice that if $i \geq m$, these are the first four summands of Line 16; if $i < m$, then we do not add $C[0]$ to the sum. In both cases, we are left with the need to subtract the sum of elements, starting from the beginning of the relevant frame, that are not a part of the last i items. Similarly to the above, we have that the sum from the beginning of the previous frame to the $i+1$ newest item is: $\sum_{d=(k-1)n+1}^{k \cdot n+m-i} x_d.$ If the last i items are all contained in the current frame (i.e., $i < m$), then we have:

$$\sum_{d=(k-1)n+1}^{k \cdot n+m-i} x_d = \sum_{d=(k-1)n+1}^{kn} x_d + \sum_{d=kn+1}^{\left\lfloor \frac{(k \cdot n+m-i)}{(\log n)^2} \right\rfloor} x_d + \sum_{d=\left\lfloor \frac{(k \cdot n+m-i)}{(\log n)^2} \right\rfloor+1}^{\left\lfloor \frac{(k \cdot n+m-i)}{\sqrt{\log n}} \right\rfloor} x_d + \sum_{d=\left\lfloor \frac{(k \cdot n+m-i)}{\sqrt{\log n}} \right\rfloor+1}^{k \cdot n+m-i} x_d.$$

In this case, we get:

1. $C[0] = \sum_{d=(k-1)n+1}^{kn} x_d.$
2. $C \left[\left\lfloor \frac{(ind-i)}{(\log n)^2} \right\rfloor \right] = \sum_{d=kn+1}^{\left\lfloor \frac{(k \cdot n+m-i)}{(\log n)^2} \right\rfloor} x_d.$
3. $SC \left[\left\lfloor \frac{(ind-i)}{\sqrt{\log n}} \right\rfloor \right] = \sum_{d=\left\lfloor \frac{(k \cdot n+m-i)}{(\log n)^2} \right\rfloor+1}^{\left\lfloor \frac{(k \cdot n+m-i)}{\sqrt{\log n}} \right\rfloor} x_d.$
4. $T[x_{\left\lfloor \frac{(ind-i)}{\sqrt{\log n}} \right\rfloor+1}, \dots, x_{ind-i}] = \sum_{d=\left\lfloor \frac{(k \cdot n+m-i)}{\sqrt{\log n}} \right\rfloor+1}^{k \cdot n+m-i} x_d.$

Here, we cancel the effect of $C[0]$ simply by not adding it as one of the summands (the *If* condition of Line 14). Quantities 2,3 and 4 are the three subtrahends of our query procedure. Finally, if $i \geq m$ we do add the value of $C[0]$, and thus in all cases we successfully compute the sum of the last i elements. ■

C Proof of Lemma 12

Proof. The algorithm utilizes three variables: \mathfrak{R} that requires $(1 + o(1)) \cdot s \log(z+1)$, \mathfrak{r} that uses $O(b \log \nu)$ bits, and o is allocated with $\lceil \log n \rceil$ bits. Overall, the number of bits used by our construction is

$$\begin{aligned} & (1 + o(1)) \cdot s \log(z+1) + O(b \log \nu) + \lceil \log n \rceil \\ & = (1 + o(1)) \cdot \lfloor n/\nu \rfloor \log(\lfloor \mu^{-1} \nu + 1 \rfloor + 1) + O(\lceil \log(n/\mu) + \log \log n \rceil \log \nu) + O(\log n). \end{aligned}$$

Since $\nu = \max\{\lfloor \mu \cdot (1 - o(1)) \rfloor, 1\}$, we get the desired bound. ■

D An $(\ell, n, 1)$ -Sliding Ranker for $\ell + 1 > 2^{\sqrt[3]{\log n}}$

Here, we detail the construction for the case of large ℓ value. We do the same splitting into frames, chunks, and sub-chunks as before. However, the large value of ℓ does not allow us to succinctly store the lookup table as before. Instead, we keep for each element the sum from the beginning of its sub-chunk, similarly to our solution in Theorem 4. Our algorithm uses the following variables:

- C - a cyclic buffer of $n/(\log n)^2$ integers, each allocated with $\lceil \log(\ell n + 1) \rceil$ bits.
- SC - a cyclic buffer of $n/\sqrt{\log n}$ integers, each allocated with $\lceil \log(\ell \log^2 n + 1) \rceil$ bits.
- $total$ - the sum of elements in the current frame.
- $subTotal$ - the sum of elements in the current sub-chunk.

Algorithm 3 An $(\ell, n, 1)$ -Sliding Ranker for $\ell + 1 > 2^{\sqrt[3]{\log n}}$

Initialization: $C \leftarrow \bar{0}, SC \leftarrow \bar{0}, ind \leftarrow 0, total \leftarrow 0, subTotal \leftarrow 0, \mathcal{W} \leftarrow \bar{0}$

- 1: **function** ADD(Element x)
- 2: $ind \leftarrow ind + 1 \bmod n$
- 3: $subTotal \leftarrow subTotal + x$
- 4: $\mathcal{W}[ind] \leftarrow subTotal$
- 5: **if** $(ind \bmod \sqrt{\log n}) = 0$ **then** ▷ End of a sub-chunk
- 6: $total \leftarrow total + subTotal$
- 7: $SC[ind/\sqrt{\log n}] \leftarrow subTotal$
- 8: $subTotal \leftarrow 0$
- 9: **if** $(ind \bmod (\log n)^2) = 0$ **then** ▷ End of a chunk
- 10: $C[ind/(\log n)^2] \leftarrow total$
- 11: **if** $(ind \bmod n) = 0$ **then** ▷ End of a frame, reset counter
- 12: $total \leftarrow 0$
- 13: **function** QUERY(i) ▷ For $i \leq n$
- 14: $addition \leftarrow 0$
- 15: **if** $i \geq ind$ **then** ▷ If not contained in current frame
- 16: $addition \leftarrow C[0]$
- 17: **return** ▷ See Figure 2

$$subTotal + SC \left\lceil \left\lfloor ind/\sqrt{\log n} \right\rfloor \right\rceil + C \left\lceil \left\lfloor ind/(\log n)^2 \right\rfloor \right\rceil + addition$$

$$- C \left\lceil \left\lfloor ((ind - i) \bmod n) / (\log n)^2 \right\rfloor \right\rceil - SC \left\lceil \left\lfloor ((ind - i) \bmod n) / \sqrt{\log n} \right\rfloor \right\rceil - \mathcal{W}[ind - i]$$

■ ind - the most recent element's index, modulo n .

■ \mathcal{W} - a cyclic array that contains for each item the sum from the beginning of its sub-chunk.

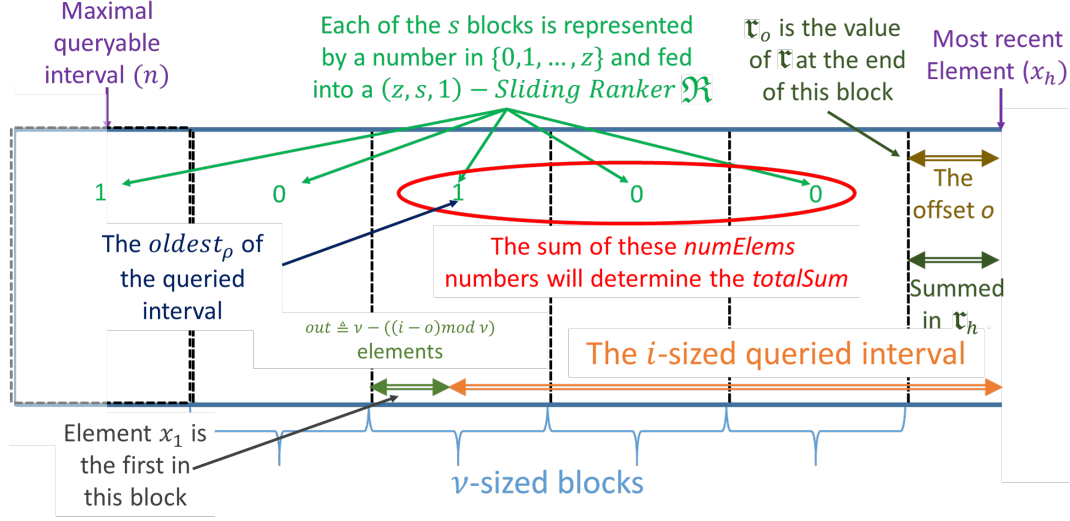
We give a pseudo code of our $(\ell, n, 1)$ -Sliding Ranker in Algorithm 3. Next, we analyze the properties of the algorithm.

► **Theorem 16.** Algorithm 3 uses $\mathcal{B}_{\ell, n, \Delta}(1 + o(1))$ memory bits for $\ell + 1 > 2^{\sqrt[3]{\log n}}$.

Proof. Similarly to the analysis in Lemma 15, the algorithm uses $o(\mathcal{B}_{\ell, n, \Delta})$ bits for keeping the chunk and sub-chunk aggregates. Here, we replaced the lookup table *and* the array of window elements by an array that stores the within-sub-chunk cumulative sum for each element. That is, each entry in the n -sized array stores a number in $\{0, 1, \dots, \ell \cdot \sqrt{\log n}\}$ and thus the array requires $n \cdot \log(\ell \cdot \sqrt{\log n} + 1) \leq n \log(\ell + 1) \left(1 + \frac{\log \log n}{\log(\ell + 1)}\right) = (1 + o(1))\mathcal{B}_{\ell, n, \Delta}$ bits overall. ■

► **Theorem 17.** Algorithm 3 is an $(\ell, n, 1)$ -Sliding Ranker.

Proof. Observe that the query procedure of our algorithm is equivalent to that of Algorithm 1, except for the part where it uses the lookup table. We now use the $subTotal$ variable to compute the sum of the current sub-chunk instead of looking it up in the table. The sum of elements that preceded the last i items in $(ind - i)$'s sub-chunk is then retrieved from $\mathcal{W}[ind - i]$. As we simply track the sum of items prior to that index in $subTotal$ and then store it in \mathcal{W} (see line 4), we get its value immediately. Thus, our estimation procedure is equivalent to that of Algorithm 1 and using Theorem 11 we establish our correctness. ■



■ **Figure 3** Theorem 14 proof's setting, with all relevant quantities that Algorithm 2 uses illustrated.

E Proof of Theorem 14

Proof. For the proof, we define a few quantities that we also use in our query procedure $numElems \triangleq \lceil \frac{i-o}{\nu} \rceil$, $totalSum \triangleq \mathfrak{R}.QUERY(numElems)$, $oldest_p \triangleq totalSum - \mathfrak{R}.QUERY(numElems - 1)$ and $out \triangleq (\nu - ((i - o) \bmod \nu))$ as in our QUERY function; see Figure 3 for illustration. We assume that the index of the most recent element is

$$h \triangleq o + out,$$

such that $o \in [\nu - 1]$ is the offset within the current block and that x_1 is the first element in the newest block of $oldest_p$. By the correctness of the \mathfrak{R} Sliding Ranker, and as illustrated in Figure 3, we have that $totalSum$ is the sum of the last $numElems$ added to \mathfrak{R} , that $oldest_p$ is the value of the element that represents the last block that overlaps with the queried window. Also notice that out is the number of elements in that block that are not a part of the window. For any $t \in \mathbb{N}$, we denote by τ_t the value of τ after the t^{th} item was added; e.g., τ_h is the value of τ at the time of the query and τ_o is its value after the last block has ended. For other variables, we consider their value at the query time.

When a block ends, we effectively perform $\tau \leftarrow \tau \bmod \tilde{\Delta}$ (lines 6 and 7) and thus:

$$0 \leq \tau_o \leq \tilde{\Delta} - 1. \quad (6)$$

Our goal is to estimate the quantity

$$S_i \triangleq \sum_{d=h-i+1}^h x_d = \sum_{d=out+1}^h x_d. \quad (7)$$

Recall that our estimation (Line 17) is:

$$\begin{aligned} \hat{S}_i &\triangleq \tau_h - \left(\tilde{\Delta} - 1/2 \right) + \tilde{\Delta} \cdot totalSum - \ell \cdot oldest_p \cdot out \\ &= \tau_{out} + \sum_{d=out+1}^h Round_b(x_d) - \left(\tilde{\Delta} - 1/2 \right) + \tilde{\Delta} \cdot totalSum - \ell \cdot oldest_p \cdot out, \end{aligned} \quad (8)$$

XX:18 Succinct Approximate Rank Queries

where the last equality follows from the fact that within a block we simply sum the rounded values (Line 4). Next, observe that we sum the rounded values in each block and that if \mathbf{r} is decreased by $k \cdot \tilde{\Delta}$ (for some $k \in \mathbb{N}$) in Line 7, then we set one of the last $numElems$ elements added to \mathfrak{R} to k . This means that:

$$\mathbf{r}_o + \sum_{d=1}^{out} Round_{\mathbf{b}}(x_d) = \mathbf{r}_{out} + \tilde{\Delta} \cdot \mathfrak{R}.QUERY(numElems) = \mathbf{r}_{out} + \tilde{\Delta} \cdot totalSum. \quad (9)$$

Plugging (9) into (8) gives us

$$\hat{S}_i = \mathbf{r}_o + \sum_{d=1}^{out} Round_{\mathbf{b}}(x_d) + \sum_{d=out+1}^h Round_{\mathbf{b}}(x_d) - \left(\tilde{\Delta} - 1/2\right) - \ell \cdot oldest_{\rho} \cdot out. \quad (10)$$

Joining (10) with (7), we can express the algorithm's error as:

$$\begin{aligned} \hat{S}_i - S_i &= \mathbf{r}_o + \sum_{d=1}^{out} Round_{\mathbf{b}}(x_d) + \sum_{d=out+1}^h \left(Round_{\mathbf{b}}(x_d) - x_d \right) - \left(\tilde{\Delta} - 1/2\right) - \ell \cdot oldest_{\rho} \cdot out \\ &= \mathbf{r}_o + \sum_{d=1}^{out} Round_{\mathbf{b}}(x_d) + \xi - \left(\tilde{\Delta} - 1/2\right) - \ell \cdot oldest_{\rho} \cdot out, \end{aligned} \quad (11)$$

where ξ is the rounding error which is defined as

$$\xi \triangleq \sum_{d=out+1}^h \left(Round_{\mathbf{b}}(x_d) - x_d \right).$$

Since each rounding of an integer $x \in \{0, 1, \dots, \ell\}$ has an error of at most $\frac{\ell}{2^{\mathbf{b}}}$, and as we round $i \leq n$ elements, we have that the rounding error satisfies

$$0 \geq \xi \geq 0 - \frac{\ell \cdot n}{2^{\mathbf{b}}} \geq -\Delta / \log n, \quad (12)$$

where the last inequality is immediate from our choice of the number of bits that is $\mathbf{b} \triangleq \lceil \log(n/\mu) + \log \log n \rceil$. We now split to cases based on the value of μ . As in the (ℓ, n, Δ) -Ranker case, we start with the simpler $\mu < 2 \cdot (1 - 1/\log n)$ case, in which $\nu = 1$ (and consequently, $out \equiv 0$). This allows us to write the algorithm's error of (11) as

$$\hat{S}_i - S_i = \mathbf{r}_o + \xi - \left(\tilde{\Delta} - 1/2\right). \quad (13)$$

We now use (6),(12) and the definition of $\tilde{\Delta}$ to obtain:

$$\hat{S}_i - S_i = \mathbf{r}_o + \xi - \left(\tilde{\Delta} - 1/2\right) \leq -1/2.$$

Similarly, we can bound it from below:

$$\hat{S}_i - S_i = \mathbf{r}_o + \xi - \left(\tilde{\Delta} - 1/2\right) \geq \xi - \left(\tilde{\Delta} - 1/2\right) \geq -\Delta + 1/2.$$

We established that if $\nu = 1$ we obtain the desired approximation. Henceforth, we focus on the case where $\mu \geq 2 \cdot (1 - 1/\log n)$, and thus $\nu = \lfloor \mu \cdot (1 - 1/\log n) \rfloor$ and $oldest_{\rho} \in \{0, 1\}$. We now consider two cases, based on the value of $oldest_{\rho}$.

1. *oldest_p* = 1 case.

In this case, we know that after the processing of element x_ν the value of \mathfrak{r} was at least $\tilde{\Delta}$ (Line 6). This implies that $\mathfrak{r}_o + \sum_{d=1}^\nu \text{Round}_b(x_d) \geq \tilde{\Delta}$ and equivalently

$$\mathfrak{r}_o + \sum_{d=1}^{out} \text{Round}_b(x_d) \geq \tilde{\Delta} - \sum_{d=out+1}^\nu \text{Round}_b(x_d).$$

Substituting this in (11), and applying (12), we get that:

$$\begin{aligned} \hat{S}_i - S_i &= \mathfrak{r}_o + \sum_{d=1}^{out} \text{Round}_b(x_d) + \xi - \left(\tilde{\Delta} - 1/2 \right) - \ell \cdot out \\ &\geq \tilde{\Delta} - \sum_{d=out+1}^\nu \text{Round}_b(x_d) + \xi - \left(\tilde{\Delta} - 1/2 \right) - \ell \cdot out \\ &\geq - \left(\sum_{d=out+1}^\nu \ell \right) + \xi + 1/2 - \ell \cdot out \\ &\geq -\Delta / \log n - \ell \lfloor \mu \cdot (1 - 1/\log n) \rfloor + 1/2 \geq -\Delta + 1/2. \end{aligned}$$

In order to bound the error from above we use (6) and (12):

$$\begin{aligned} \hat{S}_i - S_i &= \mathfrak{r}_o + \sum_{d=1}^{out} \text{Round}_b(x_d) + \xi - \left(\tilde{\Delta} - 1/2 \right) - \ell \cdot out \\ &\leq \tilde{\Delta} - 1 + \ell \cdot out - \left(\tilde{\Delta} - 1/2 \right) - \ell \cdot out \leq -1/2. \end{aligned}$$

2. *oldest_p* = 0 case.

Here, since the value of *oldest_p* was 0, we have that $\mathfrak{r}_o + \sum_{d=1}^\nu \text{Round}_b(x_d) < \tilde{\Delta}$ and thus

$$\mathfrak{r}_o + \sum_{d=1}^{out} \text{Round}_b(x_d) \leq \tilde{\Delta} - \sum_{d=out+1}^\nu \text{Round}_b(x_d) - 1.$$

We use this for the error expression of (11) to get:

$$\begin{aligned} \hat{S}_i - S_i &= \mathfrak{r}_o + \sum_{d=1}^{out} \text{Round}_b(x_d) + \xi - \left(\tilde{\Delta} - 1/2 \right) \\ &\leq \tilde{\Delta} - \sum_{d=out+1}^\nu \text{Round}_b(x_d) - 1 + \xi - \left(\tilde{\Delta} - 1/2 \right) \leq -1/2 \end{aligned}$$

We now use (6), (12), and the fact that $out \leq \nu$ to bound the error from below as follows:

$$\begin{aligned} \hat{S}_i - S_i &= \mathfrak{r}_o + \sum_{d=1}^{out} \text{Round}_b(x_d) + \xi - \left(\tilde{\Delta} - 1/2 \right) \\ &\geq \xi - \left(\tilde{\Delta} - 1/2 \right) \geq -\Delta + 1/2. \end{aligned}$$

Finally, we need to cover the case of $i \leq o$. In this case, we return $\mathfrak{r} - \left(\tilde{\Delta} - 1/2 \right)$ as the estimate. This directly follows from (6) and the fact that within a block we simply sum the rounded values (Line 4). We established that in all cases $-\Delta < \hat{S}_i - S_i < 0$, thereby proving the theorem. ■